

# Dependable Distributed Computing for the International Telecommunication Union Regional Radio Conference RRC06

Jakub T. Mościcki<sup>a</sup>, Andrea Manara<sup>b,\*</sup>, Massimo Lamanna<sup>a</sup>,  
Patricia Mendez<sup>a</sup>, Adrian Muraru<sup>a</sup>

<sup>a</sup>*CERN, Geneva, Switzerland*

<sup>b</sup>*International Telecommunication Union, Geneva, Switzerland*

---

## Abstract

The International Telecommunication Union (ITU) Regional Radio Conference (RRC06) established in 2006 a new frequency plan for the introduction of digital broadcasting in European, African, Arab, CIS countries and Iran. The preparation of the plan involved complex calculations under short deadline and required dependable and efficient computing capability. The ITU designed and deployed in-situ a dedicated PC farm, in parallel to the European Organization for Nuclear Research (CERN) which provided and supported a system based on the EGEE Grid. The planning cycle at the RRC06 required a periodic execution in the order of 200,000 short jobs, using several hundreds of CPU hours, in a period of less than 12 hours. The nature of the problem required dynamic workload-balancing and low-latency access to the computing resources. We present the strategy and key technical choices that delivered a reliable service to the RRC06.

---

## 1 Introduction

The RRC06 is the second session of the Regional Radiocommunication Conference (RRC) for the planning of the digital terrestrial broadcasting service (in band III and IV/V) in European, African, Arab, CIS countries and Iran (Fig. 1). Delegations from 104 Member States of the International Telecommunication Union (ITU [1]) gathered in Geneva to negotiate the frequency plan, from the 15th of May to the 15th of June 2006.

---

\* Corresponding author

*Email address:* `andrea.manara@itu.int` (Andrea Manara).

8 The preparation and the organization of this planning conference was man-  
9 aged by the ITU-R, the Radiocommunication Sector of the ITU. The RRC06  
10 Final Acts [2] signed by the RRC06 participants constitute a new international  
11 agreement, which comprises the new frequency plan and the procedures for  
12 its modification.

13 Analogue broadcasting has been regulated since 1961 by the Stockholm Agree-  
14 ment in Europe (ST61) and since 1989 by the Geneva Agreement for Africa  
15 (GE89). The introduction of digital technologies called for a re-planning pro-  
16 cess in order to optimize the usage of those frequency bands. The new GE06  
17 plan was designed for DVB-T (television) and T-DAB (radio) standards, but  
18 is flexible enough to accommodate future developments in digital broadcasting  
19 technologies.

20 The technical basis for this planning conference, such as the planning criteria  
21 and parameters, were established in the first session of the RRC ( RRC04  
22 [3]), which was held in Geneva in May 2004. During the RRC06 preparatory  
23 activities [4] it became evident that one component of the planning process, the  
24 compatibility analysis, was very CPU intensive. The goal of the compatibility  
25 analysis is to evaluate the interference between broadcasting requirements to  
26 identify those that can share the same channel. The analysis includes several  
27 parameters of the broadcasting requirements such as the geographic location,  
28 the signal strength and other technical characteristics.

29 The total capacity required for the compatibility analysis corresponds to sev-  
30 eral hundred CPU-days on a high-end 2006 PC. The compatibility analysis  
31 was performed in several iterations. For each iteration the RRC06 required the  
32 output of the compatibility analysis to be delivered within 12 hours. To sup-  
33 port this requirement the compatibility analysis was split in a large number  
34 of parallel calculations. The ITU-R implemented a distributed client-server  
35 infrastructure and deployed at its headquarters a dedicated farm consisting of  
36 84 high-end PCs. A distributed system based on the EGEE Grid (Enabling  
37 Grids for e-ScienE, [5]) and supported by the IT department of the European  
38 Organization for Nuclear Research (CERN) was deployed, which extended the  
39 computing capacity and improved dependability,

40 The nature of the problem required dynamic workload-balancing and low-  
41 latency access to the computing resources. This fundamental requirement was  
42 satisfied both by the ITU system, with its dedicated resources, and by the  
43 Grid system, by using high-level tools and appropriate customization of its  
44 infrastructure.

45 In this paper, we describe in section 2 the RRC06 planning process and in  
46 section 3 the computational aspects of the compatibility analysis. The imple-  
47 mentation of the ITU system is presented in section 4. The Grid-based system




Figure 1. The extent of the geographical area regulated by the GE06 Agreement.

is analyzed in section 5 and the integration of the two systems is discussed in section 6.

## 2 The RRC06 planning process

The ITU Constitution<sup>1</sup> states that “the radio-frequency spectrum is a limited natural resource that must be used rationally, efficiently and economically, in conformity with the provisions of the Radio Regulations, so that countries or groups of countries may have equitable access to it”[6].

The Radio Regulations stipulate that “Member States undertake that in assigning frequencies to stations which are capable of causing harmful interference to the services rendered by the stations of another country, such assignments are to be made in accordance with the Table of Frequency Allocations (where the frequency blocks are allocated to different radiocommunication services and to different countries) and other provisions of these Regulations”[7].

### 2.1 Frequency Planning

A frequency plan represents a key mechanism for preserving the rights of all Member States in the context of equitable access to this limited resource. Regional Radiocommunication Conferences (RRC) establish agreements concerning a particular radiocommunication service in specified frequency bands

<sup>1</sup> The ITU Constitution, the ITU Convention and the Radio Regulations are the international treaties which define the rights and obligations of ITU Member States in the domain of the international management of the frequency spectrum.

amongst participating countries. The last RRC, the RRC06, established the frequency plans (digital and analogue) for terrestrial broadcasting service (in band III and IV/V) in European, African, Arab, CIS countries and Iran. The analogue broadcasting Plan will apply only during the transition period from analogue to digital broadcasting (up to the 17 June 2015 for most Member States). After this period the broadcasting in this band will be regulated only by the digital broadcasting Plan.

Some parts of the frequency bands to be planned at the RRC06 are shared between broadcasting and other primary services (like fixed and mobile services). The planning process therefore had to take into account all services which share those bands with equal rights to operate in an interference-free environment.

## *2.2 The input data*

Member States submitted the input data to the ITU-R in the form of the so-called digital broadcasting requirements. The digital broadcasting requirements were notified as electronic files containing a set of administrative and technical parameters representing the broadcasting requirements. In addition to the digital broadcasting requirements (about 70K), the planning process had to take into account assignments to analogue television stations (about 95K) and assignments to other stations (about 10K). A fourth type of data, the so-called administrative declarations (a few million), declared that incompatibilities between digital broadcasting requirements, analogue television and other services stations may be ignored in the frequency synthesis procedure that followed the compatibilities analysis.

Radio communication services are described by administrative and technical parameters. For example, administrative parameters include the notifying administration, site name, geographic location, site altitude. Technical parameters include the power levels, assigned frequency, network topology, etc.

The digital broadcasting requirements could be submitted at the RRC06 as T-DAB (radio) or DVB-T (television) standards. Suitable data elements were provided to accommodate expected development in digital broadcasting technologies. Reference Planning Configurations served as simplified models to represent the many system variants (which differ for example in data capacity and reception modes) of the requirements. Requirements were submitted as assignments (known location and transmitter features) or as allotments (only service area known). Allotments were modeled using Reference Networks (with different number, location and power of transmitters) to approximate real networks.

104 The RRC06 planning approach was based on the protection of service areas  
105 for assignments and allotments and used the statistical model outlined in the  
106 ITU-R Recommendation P1546-1[8] to model the signal propagation.

### 107 2.3 *The planning process*

108 The ITU-R performed two planning exercises after the RRC04 and prior to  
109 the RRC06. The first planning exercise was run in June 2005 and the second  
110 in February 2006. The second planning exercise established a draft plan which  
111 served as input to the RRC06.

112 The ITU-R and the European Broadcasting Union (EBU)[9] developed the  
113 RRC06-related software. The ITU-R developed the software for data-capture,  
114 data-validation and for the display of the input data and calculation results,  
115 while the EBU developed the planning software (compatibility analysis, plan  
116 synthesis and complementary analysis). The ITU-R was also responsible for  
117 running the planning software (partly on a distributed infrastructure), pro-  
118 ducing and delivering results in due time.

119 At the RRC06 the frequency plan was established in an iterative way, as  
120 outlined in Fig.2 The delegations engaged in bilateral and multilateral coordi-  
121 nation and negotiation efforts which resulted in a new set of refined digital  
122 broadcasting requirements at the end of every week. Over the weekends the  
123 ITU-R performed the validation of the data and the compatibility analysis  
124 and synthesis calculations. The output of these calculations and the refined  
125 frequency plan were the input for the negotiations in subsequent week, with  
126 the last (fourth) iteration constituting the basis for the final frequency plan.  
127 In order to assist groups of negotiating Member States, partial calculations  
128 were performed for parts of the planning area in between two global iterations.

129 The compatibility analysis consisted of the calculation of the interference be-  
130 tween digital broadcasting requirements and other primary services stations.  
131 For each requirement the compatibility assessment produces a list of incompat-  
132 ible requirements and a list of available channels. Three types of compatibility  
133 analyses were needed, for both UHF and VHF frequency bands: digital versus  
134 digital (d2dUHF and d2dVHF), digital versus other services (d2oUHF and  
135 d2oVHF) and other services versus digital (o2dUHF and o2dVHF).

136 These lists were the input to the plan synthesis process, which determined a  
137 suitable frequency for each requirement in order to avoid harmful interference  
138 and to maximize the number of requirements satisfied. The RRC06 decided  
139 to protect analogue broadcasting services during the implementation of the  
140 digital broadcasting requirements rather than during the establishment of the  
141 plan to maximize the number of requirements satisfied. For this reason each it-

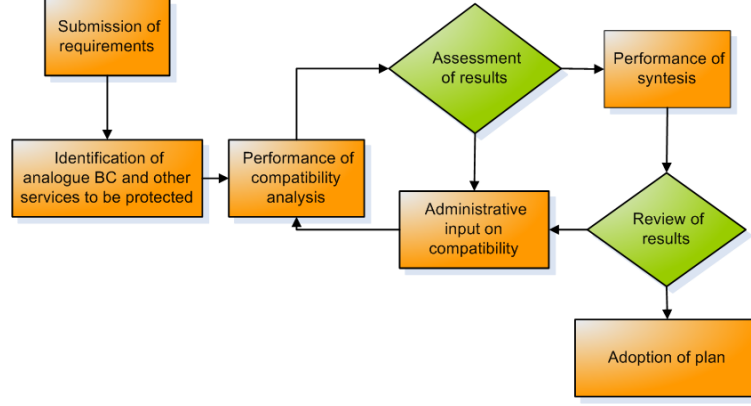


Figure 2. ITU negotiation workflow.

eration included a complementary analysis, which determined which analogue television assignments may suffer interference from the implementation of a given digital broadcasting assignment or allotment.

During pre-conference preparatory planning activities only 34% of requirements were satisfied. For the first iteration of the RRC06 the percentage increased to 64% (UHF) and 74% (VHF), to reach a satisfactory 93% (UHF) and 98% (VHF) for the final plan.

### 3 The computational challenge

The compatibility assessment is CPU-intensive. In the compatibility analyses each requirement must be run against all the others, for six different types of analysis (d2dUHF, d2dVHF, d2oUHF, d2oVHF, o2dUHF, o2dVHF). In this paper we use the term *atomic calculations* to refer to individual, indivisible calculations defined in compatibility analysis datasets. The term *task* refers a unit of work which corresponds to a set of atomic calculations. The term *job* is used in the context of Grid job submission only.

For the first planning exercise the atomic calculations were clustered in tasks of 100 for all types of analyses. With the limited resources available at that time, that exercise took about one week (elapsed time), for an integrated 90 CPU days.

The detailed study revealed an exponential distribution of the requirement processing time which spans almost three orders of magnitude (Fig. 3). The huge variation in running time depends, among other parameters, on the number of acceptable channels specified in the digital broadcasting requirement, the requirement type (assignment versus allotment), the network topology and signal propagation zones specific to the geographical area of the Member




Figure 3. Distribution of the number of processed requirements per hour for the d2dUHF analysis as a function of the Member State. Data for the first planning exercise.

iteration	d2dUHF	d2dVHF	d2oUHF	d2oVHF	o2dUHF	o2dVHF
1	3(3)	5(5)	100(100)	100(100)	100(100)	100(100)
2	4(3)	4(10)	50(100)	50(100)	100(100)	100(100)
3	2(3)	2(5)	50(100)	50(100)	50(100)	50(100)
4	2(3)	2(10)	50(100)	50(100)	50(100)	50(100)

Table 1

Compatibility analysis granularity for the RRC06 iterations for Grid and ITU (in parenthesis) system.

State.

Further investigation showed that a complete static optimization of the load <sup>2</sup> was not possible due to the unpredictable nature of the data as the Member States could change their requirements before each RRC06 iteration. On the other hand, there was clearly a need to create smaller clusters for the most CPU demanding type of analysis d2dUHF and d2dVHF, minimizing the spread between the shortest and longest tasks. Table 1 shows the granularity chosen for the different types of analysis in the RRC06 iterations for the Grid and ITU systems. The granularity was adjusted manually in between the iterations. The load balancing was handled dynamically at runtime.

The workload for each compatibility analysis run at the RRC06 corresponded

<sup>2</sup> The static optimization of the load is an ability to a priori cluster the requirements, so that the execution time of each cluster is equal.




Figure 4. Architecture of the ITU dedicated system.

178 to some several hundred CPU hours. Additionally the workload was to be  
 179 completed within a deadline of a few hours. The time constraints were critical:  
 180 an hypothetical problem with timely delivery of analysis results could have  
 181 resulted in a failure of international negotiations.

182 The total CPU demand decreased with each RRC06 iteration. Member States  
 183 decreased the number of requirements and the number of acceptable channels  
 184 for each requirement, reducing therefore the total workload at each analysis  
 185 iteration. Finally, as the frequency plan was refined during successful negotia-  
 186 tions between the Member States, the number of conflicting requirements also  
 187 decreased. The CPU demands for the ITU and Grid systems is presented in  
 188 the next sections.

## 189 4 ITU system

190 The ITU system consisted of a client-server distributed system running on a  
 191 dedicated PC farm. The farm resources evolved in time. Initially it consisted of  
 192 six high-end dedicated PCs complemented by some tens of ITU staff desktop  
 193 PCs, available only overnight and during weekends. Using this configuration,  
 194 the calculations for the first planning exercise required about one week, show-  
 195 ing that the running time was an outstanding issue in preparation for the  
 196 RRC06. The ITU-R therefore decided to buy a PC farm, which was deployed  
 197 within ITU headquarters by the ITU Infrastructure Services department (ITU  
 198 IS). In its final configuration at the RRC06 the farm was composed of 84 high-  
 199 end dedicated 3.6 GHz hyper threading PCs. Accurate measurements showed  
 200 that hyper threading permits to gain about 30% in computing time by running  
 201 two tasks in parallel on one PC with respect to the situation when the same  
 202 tasks are run sequentially.

203 To cope with redundancy and logistic issues (available space, power and cool-



Iteration	$N_{calc}$	$N_{task}$	$t_{total}$	$t_{clients}$
1	173K	26K	5.9h	621h
2	168K	23K	4.1h	463h
3	154K	23K	3.4h	300h
4	155K	21K	2.6h	205h

Table 2

Performance of the ITU system (84\*2 simultaneous processes) during compatibility analysis calculations




Figure 5. Distribution of the elapsed time for the ITU system during RRC06 iteration 2.

ing consideration), ITU-IS decided to deploy the farm into two separate clusters. The first cluster consisted of 47 PCs and was equipped with optical fibers and a 1Gb/s network switch, while the second cluster consisted of 37 PCs with a slower 200Mb/s network switch. This configuration did not significantly impact on the performance of the system.

The architecture layout is presented in Fig. 4. The system was implemented with Perl scripts installed as Windows services and a custom communication protocol based on UDP/IP. The UDP packets carried information on the executable to be run and on the relevant input parameters. In the reliable internal network of the ITU farm the packet loss was not a problem. The server implemented two Windows services, a Listener and a Dispatcher, responsible for task submission, task management and workload balancing. To cope with high-load, the TaskQueue file ensured asynchronous operation of the system and prevented packet lost. The system automatically managed the task status and resubmitted the ones which were not completed.




Figure 6. Number of running processes as a function of time during RRC06 iteration 2.

219 The clients implemented two Windows services, the TaskManager responsible  
 220 for running tasks according to Dispatcher requests and the TaskController  
 221 responsible for monitoring and control operations. A web application (imple-  
 222 mented with ASP.NET and C#) running on a dedicated machine (WebInter-  
 223 face), provided monitoring and control interfaces to operate the system.

224 In the first phase, the client installation on non-dedicated resources (desktop  
 225 PCs) was implemented using a MSI-compatible installation procedure man-  
 226 aged by Windows Systems Management Server (SMS). In the dedicated farm,  
 227 the software and data were deployed on a shared folder and copied directly  
 228 to the client PCs. MD5 checksums were performed to insure data consistency.  
 229 At system startup the server automatically triggered the software and data  
 230 installation at the client.

231 The system supported 2\*84 simultaneous tasks most of the time with negligible  
 232 job loss. Software and data installation involved 350 MB to be deployed in  
 233 2\*84 folders and took on average 15 minutes for the entire farm.

234 The performance of the ITU system is reported in Table 2, where the total  
 235 workload of atomic calculations  $N_{calc}$ , the number of tasks  $N_{task}$ , the total time  
 236 to complete the iteration  $t_{total}$  and the integrated elapsed time on the clients  
 237  $t_{clients}$  are shown for each iteration. The distribution of the tasks processing  
 238 time for the ITU system during iteration 2 of the RRC06 is shown in Fig. 5.  
 239 The evolution of the number of running processes as a function of time dur-  
 240 ing RRC06 iteration is shown in Fig. 6. This last figure illustrates interesting  
 241 features of the ITU system: the dynamic load balancing (about 96% of the  
 242 clients complete processing tasks practically at the same time) and limited

243 submission latency (about 15 minutes, the time necessary for the clients to  
244 download the latest version of software and data at server start-up).

245 Taking into consideration also the four runs of complementary analysis and the  
246 partial runs during multilateral negotiations, the ITU system at the RRC06  
247 ran more than 180 thousand tasks for an overall integrated elapsed time of  
248 4500 CPU/hours, i.e. more than half a CPU year.

## 249 5 Grid system

250 Enabling Grids for E-science (EGEE) is a globally distributed system for  
251 large-scale batch job processing. At present it consists of around 300 sites in 50  
252 countries and offers more than 80 thousand CPU cores and 20 PB of storage  
253 to 10 thousand users around the globe. EGEE is a multidisciplinary Grid,  
254 supporting users in both academia and business, in many areas of physics,  
255 biomedical applications, theoretical fundamental research and earth sciences.  
256 The largest user communities come from the High-Energy Physics, and in  
257 particular the experiments active at the CERN Large Hadron Collider (LHC).

258 The EGEE Grid has been designed and operated for non-interactive processing  
259 of very long jobs. A set of complex middleware services integrate computing  
260 farms and the batch queues into a single, globally distributed system. The ac-  
261 cess to the distributed resources is typically controlled by the fair-share mech-  
262 anisms, ensuring usage of resources by groups of users according to predefined  
263 policies. In typical configurations a large number of users share individual  
264 computing resources across multiple Virtual Organizations (VOs)<sup>3</sup> This ar-  
265 chitecture is suitable for high-throughput computing but is not efficient for  
266 high-performance, short-deadline, dependable computing which is stipulated  
267 by the RRC06 compatibility analysis application.

268 In the EGEE Grid environment and on a short time-scale these requirements  
269 may only be implemented if high-level tools are used to control the job work-  
270 load and the Grid infrastructure is appropriately customized.

### 271 5.1 *The tools*

272 To run RRC06 compatibility analysis application Ganga and DIANE tools  
273 were used.

---

<sup>3</sup> Virtual Organization is a group of users sharing the same resources. Members of one Virtual Organization may belong to different institutions.

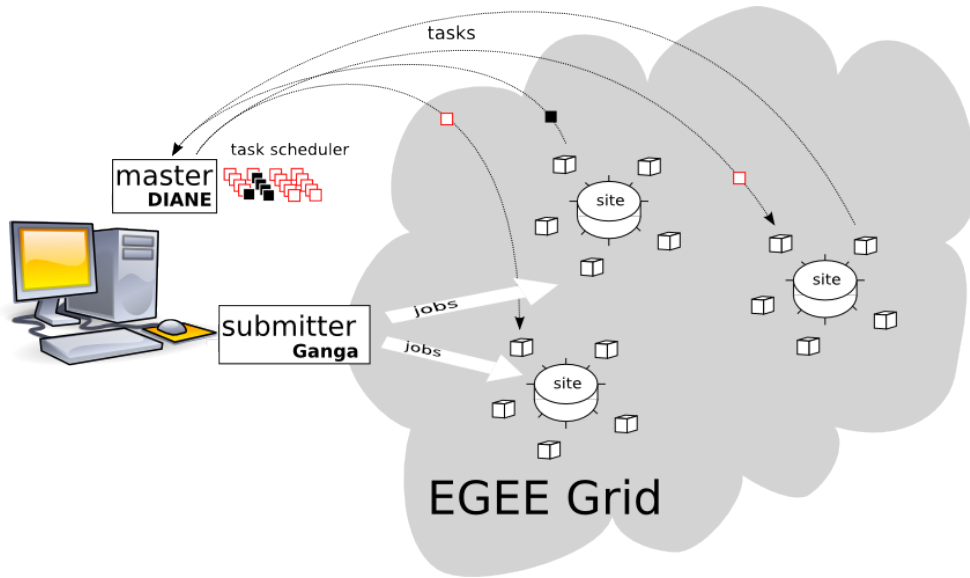


Figure 7. Overview of the Grid system based on Ganga/DIANE.

274 Ganga provides a uniform and flexible interface to submit, track and manip-  
 275 ulate jobs [18]. DIANE is an agent-based job scheduler which provides fault-  
 276 tolerant execution of jobs, dynamic workload-balancing and reduced overhead  
 277 in accessing the computational resources [19].

278 The outline of the architecture is presented in Fig.7. Worker agents are sub-  
 279 mitted to the Grid and pull the tasks from the Master server which controls  
 280 the distribution of the workload. The system is fault-tolerant and may run  
 281 autonomously: a Worker agent which fails to complete the assigned calcula-  
 282 tions is replaced by another Worker agent. The overhead of scheduling the  
 283 calculations is negligible in comparison with the overhead of classic Grid job  
 284 submission. The system dynamically reacts to changing workload and provides  
 285 dynamic load-balancing. The results of the compatibility analysis of the re-  
 286 quirements are directly uploaded to the Master server. The implementation of  
 287 the RRC06 system on the EGEE Grid was based on DIANE 1.5.0 and Ganga  
 288 4.1.

289 The input data, including the specification of the digital broadcasting require-  
 290 ments and the tuned compatibility analysis application, were distributed to the  
 291 collaborating Grid sites shortly before the analysis was launched. The 100MB  
 292 installation package was deployed into the directory mounted on a shared file  
 293 system accessible by all worker nodes of a collaborating Grid site (so called  
 294 “software areas”). The installation was managed by separate grid jobs running  
 295 with the credentials of the VO manager and using MD5 check-sums to assure  
 296 consistency of the installation tarballs. The installation was automated and  
 297 the installation jobs checked periodically to download the installation pack-  
 298 ages available in a central repository at CERN. This allowed to automatically  
 299 distribute the new installation packages in 15 minutes after the ITU-R made

iteration	$N_{calc}$	$N_{task}$	$t_{total}$	$t_{worker}$	$N_{worker}$	$r_{fail}$
1	243K	26K	6h40m	425h	190	$<3\text{e}-4$
2	237K	23K	6h30m	332h	125	$4\text{e}-5$
3	224K	40K	1h35m	192h	210	0
4	218K	39K	1h5m	151h	320	0

Table 3

Summary of RRC06 compatibility analysis iterations.

300 them available.

301 The ITU personnel updated the software packages with 2 hours' notice. In this  
302 time window the grid system had to be up and ready to start the computation  
303 at full speed, as soon as the update was available.

## 304 5.2 The infrastructure

305 The access to the computing resources on the Grid for the RRC06 use was im-  
306 plemented using the GEAR Virtual Organization (vo.gear.cern.ch). The CPU  
307 demand for RRC06 was much smaller than typical Grid applications which  
308 require huge throughput over very long periods of time. However, conversely  
309 to many other Grid applications, availability of resources within well-defined  
310 and strict time constraints was critical. Therefore a number of high-availability  
311 centres in the EGEE Grid <sup>4</sup> were involved. The resources at these centres were  
312 not dedicated to the RRC06 activity, however the job priority parameters were  
313 adjusted during short periods of intensive processing of the RRC06 compat-  
314 ibility analysis (the weekends between the major conference iterations). On  
315 average 300 CPUs were observed to be available at all times with occasional  
316 peaks of c.a. 600 CPUs.

317 Redundant deployment of key services, such as the Master servers, Grid User  
318 Interfaces and Resource Brokers [15] allowed for fail-over in case of problems.  
319 For storing the application output the AFS and local filesystem were used  
320 simultaneously.

## 321 5.3 Analysis of the system

322 The summary of RRC06 iterations is presented in Table 3. For each anal-  
323 ysis iteration the total workload consisted of  $N_{calc}$  atomic calculations. The

<sup>4</sup> CERN, CNAF+few other sites(I), PIC(E), DESY(D), MSU(RU) ,  
CYFRONET(PL)




Figure 8. Run 3 workload. Resolution=60s.

324 calculations were executed in bunches according to previously defined static  
 325 clustering (section 3). The  $N_{task}$  tasks were distributed dynamically to the  
 326  $N_{worker}$  Worker agents. The Worker agents were submitted as *jobs* and exe-  
 327 cuted on the Grid worker nodes.  $t_{total}$  is the makespan or the total time to  
 328 complete the compatibility analysis.  $t_{worker}$  is the integrated elapsed time on  
 329 the worker nodes.  $r_{fail}$  is the reliability of the system and corresponds to the  
 330 number of failed tasks which could not automatically recover. With fewer than  
 331 10 lost tasks in run 1 and one lost task in run 2 the reliability of the system  
 332 exceeded by few orders of magnitude the reliability of the Grid infrastructure.

333 Contrary to the ITU system which used a fixed set of resources, in the Grid  
 334 resources are dynamic: a different set of worker nodes is used at each iteration.  
 335 The worker node characteristics such as the CPU and memory also show  
 336 large variations. Therefore a direct comparison of  $t_{total}$  and  $t_{worker}$  parameters  
 337 between ITU and Grid runs is not possible.

338 The efficiency of the system depends on the Grid job submission latency, effi-  
 339 ciency of task scheduling and workload balancing. Fig. 8,9 show the workload  
 340 distribution for selected runs.  $N_w$  worker agents are submitted at  $t_0 = 0$ . In  
 341 the submission phase,  $t < t_1$ , the throughput of the system is limited by the  
 342 submission latency. As the pool of worker nodes increases the target of  $N_w$   
 343 workers is reached at time  $t_1$ . In the main processing phase,  $t_1 < t < t_2$ ,  
 344 the pool of worker nodes remains stable and the system throughput mainly  
 345 depends on the efficiency of scheduling. At time  $t_2$  the number of remaining  
 346 tasks becomes smaller than the number of processors in the pool. In this phase  
 347 the execution time is dominated by the workload-balancing effects from few  
 348 slowest tasks.




Figure 9. Run 4 workload. Resolution=60s. The point  $t_1$  was selected arbitrarily. In run 4 two parallel master servers were used and this figure corresponds to one of the masters and half of the total workload.




Figure 10. Run 3 profile.




Figure 11. Run 4 profile.

349 The number of available worker nodes may vary significantly in the Grid from  
 350 one run to another. The contribution of the job submission latency to the total  
 351 execution time may be approximated by the area between the target line and  
 352 the worker pool size curve. In run 3 the latency of job submission corresponded  
 353 to 12% of the total execution time, whereas in run 4 it corresponded to 48%:  
 354 33% in the submission phase and 15% in the main processing phase.

355 The integrated difference between the worker pool size and the number of  
 356 busy workers corresponds to the scheduling overhead. This overhead includes  
 357 the network latency and throughput as well as the task handling efficiency  
 358 of the master server. In run 3 the scheduling overhead in the submission and  
 359 processing phases corresponded to 2-3%. In run 4 the 30% scheduling overhead  
 360 in the submission phase was observed and 10% in the processing phase.

361 The unbalanced execution of the slowest tasks in the last phase contributes  
362 26% of the total execution time in run 3 and to 5% in run 4. In this phase the  
363 utilization of available resources was very low, 5% in run 3 and 20% in run 4.  
364 The majority of the workers in the pool remained idle while the few remaining  
365 tasks were being finished.

366 The striking difference of scheduling and workload-balancing efficiency be-  
367 tween runs 3 and 4 may be explained by the task scheduling order which  
368 reflects the internal input data structure. The run profile plots are shown in  
369 Fig. 10, 11. Point (t,w) in the run profile represents a task completed by worker  
370 w at time t. In run 4 the tasks are drawn directly from the input data in the  
371 natural order and clusters of very short tasks created a very high load on the  
372 server. The long tasks were processed in the middle of the run and did not  
373 affect the overall load-balancing. In run 3 the tasks were selected in a random  
374 order by the scheduler. The momentary load on the server was reduced. The  
375 tasks were scheduled more uniformly across the entire run. There were a few  
376 long tasks at the end of the run that resulted in poor load-balancing.

377 The intrinsic job submission latency in the Grid prevents the running of a large  
378 number of short jobs in a short time, unless user-level tools such as DIANE  
379 are used. For RRC06 using DIANE allowed to reduce the Grid overheads  
380 and provided efficient management of a large number of tasks. Additionally a  
381 runtime workload balancing allowed to evenly distribute a workload without  
382 precise, a priori knowledge of the task execution times in the dataset. The  
383 overhead reduction and workload balancing were the crucial factors of the  
384 successful usage of the Grid for the RRC06.

## 385 6 System Integration

386 The Grid and ITU systems were integrated at the monitoring level using the  
387 MonALISA framework (Monitoring Agents in A Large Integrated Services  
388 Architecture, developed by Caltech University [20]). MonALISA provides a  
389 set of pluggable distributed services for monitoring, control, management and  
390 global optimization for large scale distributed systems.

391 To collect and combine monitoring information from both ITU and Grid sys-  
392 tems, the following software components were deployed: instances of MonAL-  
393 ISA collector service, web-enabled data visualization repository and custom  
394 ApMon monitoring sensors on worker nodes (Fig. 12). ApMon, the monitoring  
395 API, allows to send fine-grained custom monitoring parameters into the Mon-  
396 ALISA collector service. The ApMon uses UDP datagrams to transport the  
397 XDR-encoded information [21] and includes a sequence number to verify the  
398 integrity of all monitoring reports. In addition, ApMon provides out-of-the-






Figure 12. System integration via Mona Lisa monitoring.

box system monitoring of the host, including usage of system resources such as memory or CPU. Monitoring parameters of ApMon, such as monitoring frequency and collector destination, may be dynamically configured by remote services. ApMon implementations are provided for different programming languages, including C, C++, Java, Perl and Python. The cross-language support has proven to be useful in the case of RRC06 as the ITU system was built in Perl while the Grid used Python.

Using pluggable modules, the MonALISA collector has been customized to aggregate fine-grained data from Grid worker nodes and ITU farm nodes to produce in real-time, higher level reports and charts. Fig.13 shows the total workload executed by ITU clusters and the EGEE sites. The ITU clusters are reported as `RRC06-1.itu.org` and `RRC06-2.itu.org`.

The complementary usage of Grid Unix-based and Windows-based resources for numerical computations, required compilation of application software on both platforms and verification of output in terms of numerical accuracy.

## 7 Conclusions and Outlook

The dual system presented in this paper contributed to the success of the RRC06 Conference which resulted in a new international treaty.

Seamless access to resources from Grid and corporate infrastructures demonstrated in this paper may be beneficial for other user communities. A typical use-case could include dedicated in-situ resources for fast response and Grid resources when facing peak demand. In such a scenario the Grid could provide a competitive alternative to traditional procurement of resources. At RRC06 the Grid delivered dependable peak capacity to an organization which normally does not require a large permanent computing infrastructure. The Grid




Figure 13. Total workload executed in Grid and ITU clusters.

424 was successfully used in a new area to provide a dependable just-in-time ser-  
 425 vice. ITU personnel needed limited support and training to adopt the Grid  
 426 technology for RRC06. This demonstrates the maturity of Grid technology for  
 427 usage in new scientific communities and technical activities.

428 The outcome of RRC06 was the GE06 frequency plan which is a part of  
 429 an international agreement. Modifications to the GE06 Plan may require a  
 430 coordination examination to determine Member States potentially affected. To  
 431 bring into use a new broadcasting station a conformity examination is required  
 432 to verify that the proposed implementation does not cause more interference  
 433 than foreseen by the GE06 Plan. Both examinations may require intensive  
 434 calculations. In addition, some Member States have already expressed the  
 435 possible need for re-planning parts of the GE06 planned bands, a process  
 436 which would imply a similar (smaller scale) approach to the one adopted at  
 437 the RRC06.

438 In order to prepare for future events which may require even more comput-  
 439 ing capabilities than the RRC06, paradigms such as Cloud computing could  
 440 be investigated, where dynamically scalable resources are provided as a ser-  
 441 vice over the Internet. A system integrating local, grid and cloud resources  
 442 would allow Member States to submit via an existing ITU web portal time-  
 443 consuming calculation requests and, at the back-end, to schedule and execute  
 444 jobs transparently on the integrated infrastructure. Such a pilot project could  
 445 be a continuation of the system accomplished for the RRC06 and a potential  
 446 area of future collaboration between ITU and CERN.

## 447 8 Acknowledgments

448 The authors would like first of all to thank the sites pledging resourcing for this  
449 activity. The sites are: CNAF (Bologna, Italy) complemented by a few other  
450 sites belonging to the Italian infrastructure (Grid IT), CYFRONET (Krakow,  
451 Poland) DESY (Hamburg and Zeuthen, Germany), MSU (MOSCOW, Russia),  
452 PIC (Barcelona, Spain). Their willingness to share their resources was an  
453 essential condition for the success of this activity.

454 We would like to thank the MonALISA team (CERN/Caltech) and in par-  
455 ticular Iosif Legrand for support. We would also like to thank the ITU-IS  
456 for the ITU farm set-up and continuous support and all the ITU-R staff in-  
457 volved in the data preparation for the RRC06. We are grateful to M. Cosic  
458 and D.Botha for providing the compilation of the compatibility analysis ex-  
459 ecutable for Linux and Windows platform, and to J. Boursy for encouraging  
460 the developing of an ITU distributed computing system several years before  
461 it was needed at the RRC06. A special thank to P.N. Hai (ITU), who coor-  
462 dinated the overall processing at the RRC06, for encouragement, vision and  
463 useful feedback. Finally a special thank goes the management of the ITU-R  
464 and CERN IT department, in particular to T. Gavrilov and W. von Rden  
465 (CERN IT Dept head) for encouragement and useful discussions.

## 466 References

- 467 [1] International Telecommunication Union, <http://www.itu.int>
- 468 [2] Final Acts of the Regional Radiocommunication Conference for planning  
469 of the digital terrestrial broadcasting service in parts of Regions 1 and  
470 3, in the frequency bands 174-230 MHz and 470-862 MHz (RRC-06),  
471 <http://www.itu.int/ITU-R/conferences/rrc/rrc-06/index.asp>
- 472 [3] Regional Radiocommunication Conference 2004 for planning of the digital  
473 terrestrial broadcasting service in parts of Regions 1 and 3, in the frequency  
474 bands 174-230 MHz and 470-862 MHz (RRC-04),(Geneva, Switzerland, 10 May-  
475 28 May 2004), <http://www.itu.int/ITU-R/conferences/rrc/rrc-04/index.asp>
- 476 [4] Intersessional Activities (RRC-  
477 04/06),(Geneva, Switzerland, May 2004-May 2006), <http://www.itu.int/ITU-R/conferences/rrc/rrc-04/intersession/index.asp>
- 479 [5] Enabling Grid for E-science (EGEE) home page: <http://www.eu-egee.org>
- 480 [6] Constitution of the ITU, Chapter VII, Art. 44, "Use of the Radio-Frequency  
481 Spectrum and of the Geostationary-Satellite and Other Satellite Orbits",  
482 <http://www.itu.int/net/about/basic-texts/constitution/chaptervii.aspx>

- 483 [7] International Telecommunication Union, Radio Regulations, Article 4.1
- 484 [8] Recommendation ITU-R P.1546-3 Method for point-to-area predictions for  
485 terrestrial services in the frequency range 30 MHz to 3000 MHz
- 486 [9] European Broadcasting Union home page, <http://www.ebu.ch>
- 487 [10] General updated information on the LHC programme can be found on the  
488 CERN web site (<http://www.cern.ch>). A recent review article on the first 2  
489 years of LHC is: Fabiola Gianotti “Physics during the first two years of the  
490 LHC”, New J. Phys. 9 (2007) 332. DOI: 10.1088/1367-2630/9/9/332
- 491 [11] Open Science Grid (OSG) Web Page, <http://www.opensciencegrid.org>
- 492 [12] Nordic Data Grid Facility (NDGF) Web Page, <http://www.ndgf.org>
- 493 [13] Ian Foster and Carl Kesselman, “The GRID: Blueprint for a New Computing  
494 Infrastructure”, Morgan Kaufmann, 1998
- 495 [14] S. Bethke et al., “Report of the Steering Group of the LHC Computing Review”,  
496 CERN/LHC/2001-004, CERN/RRB-D 2001-3, 22 February 2001
- 497 [15] LHC Computing Grid (LCG) home page: <http://cern.ch/lcg>
- 498 [16] Models of Networked Analysis at Regional Centers for LHC Experiments  
499 (MONARC) project home page, <http://cern.ch/monarc>
- 500 [17] The LCG Editorial Board, “LHC Computing Grid Technical Design Report”,  
501 LCG-TDR-001, CERN-LHCC-2005-024, June 2005
- 502 [18] J.T.Mościcki et.al.: Ganga: a tool for computational-task management and easy  
503 access to Grid resources; arXiv:0902.2685v1
- 504 [19] Distributed Analysis Environment, <http://cern.ch/diane>
- 505 [20] Monitoring Agents Using a Large Integrated Services (MonALISA) project home  
506 page: <http://monalisa.cern.ch/monalisa.html>
- 507 [21] RFC1014 - XDR: External Data Representation standard,  
508 <http://www.faqs.org/rfcs/rfc1014.html>
- 509 [22] C. Germain-Renaud, C. Loomis , J. T. Mościcki and R. Texier, *Scheduling for*  
510 *Responsive Grids*, *J. Grid Computing* 6, (2008) 15-27